Adding Semantic Information into Data Models by Learning Domain Expertise from User Interaction

Nathan O. Hodas, Alex Endert

Abstract— Interactive visual analytic systems enable users to discover insights from complex data. Users can express and test hypotheses via user interaction, leveraging their domain expertise and prior knowledge to guide and steer the analytic models in the system. For example, semantic interaction techniques enable systems to learn from the user's interactions and steer the underlying analytic models based on the user's analytical reasoning. However, an open challenge is how to not only steer models based on the dimensions or features of the data, but how to add dimensions or attributes to the data based on the domain expertise of the user. In this paper, we present a technique for inferring and appending dimensions onto the dataset based on the prior expertise of the user expressed via user interactions. Our technique enables users to directly manipulate a spatial organization of data, from which both the dimensions of the data are weighted, and also dimensions created to represent the prior knowledge the user brings to the system. We describe this technique and demonstrate its utility via a use case.

Index Terms—Sense-making, Feature Reduction, Information Theory, Knowledge Generation

arXiv:1604.02935v1 [cs.HC] 6 Apr 2016

1 Introduction

Visual analytic systems facilitate a cognitive process of discovery by combining data analytic and information visualization capabilities . These technologies leverage computational models to generate visual representations of data. Through user interaction, analysts make use of analytical reasoning to pose questions, form hypotheses, and gain insights into the data. For exploratory data analysis, this process can be referred to as sensemaking [5].

The analytic models used in visual analytics typically approximate some characteristics, metrics, attributes, or other features about the data. Thus, user interactions are typically designed to augment the analytic or visual parameters. More recently, semantic interaction techniques enable analytic model steering via inferring data characteristics that are of interest from the user based on their exploratory user interactions [4]. However, one limitation of such approaches is that they attempt to fit the inference of the user's feedback into the existing feature space in the data model (e.g., keywords for text corpora). Analysts create mental constructs and concepts about the data which do not necessarily reflect in the data features, but extend into the broader context and features existent in the domain of the data being explored. Thus, *how can visual analytic systems infer such additional, semantic dimensions about the data from the user?*

In this paper, we demonstrate a system for allowing a user to create additional semantic features and leverage these features for image classification and accelerating the sense-making process. The system presents a canvas filled with randomly arranged images to the user, and the user positions images according to their internal mental model, leveraging their previous knowledge and task-specific motivations. We call this system ActiveCanvas, because it does not simply allow the user to shuffle images around; it actively participates in the sensemaking loop by using information theoretic tools to extract relevant features from the images and then repositioning the images in an attempt to better reflect the user's mental model. The information theoretic tools implicitly capture the analytic reasoning of the user, and it allows the system to build knowledge over time by recording the position of users and saving them as new features to be utilized in the future.

- *Nathan Hodas is with Pacific Northwest National Laboratory. E-mail: nathan.hodas@pnnl.gov.*
- *Alex Endert is with Georgia Tech. E-mail: endert@gatech.edu.*

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014; date of publication xx xxx 2014; date of current version xx xxx 2014.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

2 Related Work

Spatial workspaces are effective visual metaphors for supporting sensemaking and discovery tasks. For example, Andrews et al. showed how the spatial positioning, grouping, and re-grouping of documents was an effective method to help analysts conduct a sensemaking task [1]. Their work found that users were able to leverage the flexibility of a large, high-resolution to organize their information. Throughout the process, the analysts made use of spatial constructs including lists, piles, sequences, and categorizations of the information. Over time, these spatial constructs evolved as the analysts gained more insights about the data. This process of incrementally refining a spatial layout based on one's own increased understanding of the data is called *incremental formalism* [6].

User interaction is critical to the success of sensemaking and discovery tasks using visual analytics. Prior work has analyzed user interaction as a data source to understand more about the user, the domain, and the data being analyzed. For example, Dou et al. showed how the analysis of user interaction logs can uncover low-level strategies and tasks of analysts performing financial analysis tasks [3]. Similarly, Brown et al. analyzed user interactions for a simple search task (Where's Waldo) to discover patterns of interactions that could accurately predict fast and slow task completion times [2].

3 ActiveCanvas

The ActiveCanvas is designed to enable the user to leverage the hard work of other users to accelerate the work of a new user sorting images on a 2d virtual space. Imagine that a user has a latent mental model of the data they are trying to organize, denoted by X_L. X_L represents the implicit features a user internally leverages to understand the images in front of them. Initially, the system would present them with naive arrangements of images, each of which the user would arrange on the screen to create positions Y. Thus, for this limited set, Y explains X_L, to the best of the user's intentions. A key assumption is that the user places items close together on the screen (Y-space) because they are similar in the latent space X_L. Thus, if we can find a way to slightly alter the arrangement of items on the screen to maximize mutual information with the feature vector provided by the system, these new positions will likely have high mutual information with the user's mental model.

To this end, when a user has arranged some initial images to their satisfaction, the ActiveCanvas attempts to slightly alter the new positions of those images, Y' such that $MI(X : Y')$ increases, where MI is the mutual information. We then predict the position of the remaining items using the new positions of the touched items as a training set. using RBF SVM. If the user is not satisfied, the user continues to

adjust images or move additional images to produce better estimates of mutual information and more training data for the SVM. After the user has become satisfied with the position of all items on the screen, they may then commit the positions as two new features, changing $X \rightarrow [XY_{final}]$. This enables other users to leverage the hard work of the first uers.

This robustness to imprecision allows a lazy user to benefit from the diligent work of more dedicated users. If two users have classified images according to two different strategies, the lazy user need not position their images with as much rigor as would otherwise be required to distinguish clear clusters. Upon sending the system the lazily-positioned images, ActiveCanvas leverages the features (e.g., positions chosen and recorded by the other users) to infer the intended clusters of the lazy user. After being satisfied, each user commits their arrangement back to the system, creating a new features set comprised of the original feature vector plus the (x,y) positions from each user.

To exemplify this, consider the following sample workflow, worked through in Fig. 1. As above, the user approaches ActiveCanvas to make sense of images chosen because they contain one of {frog,dice,deer,elephant,pram}, although they may be interested in the images for any reason. Previous users have utilized ActiveCanvas on the same dataset to make clusters based on their personal desired sorting, including {human, not human}, {light background, dark background}, {outdoors, not outdoors}. {animal, not animal}, as well as according to the base classes. After sorting the images according to their preference, they each save the positions of the images as new features. The new user approaches the canvas initially seeing a random arrangement of pictures. They then move only 8 images into rough piles, or even singleton piles, shown in Fig. 1(a). The user then asks ActiveCanvas to refine the arrangement according to the implicit mental model ActiveCanvas uncovers from the 8 images the user touched. This results in an initial arrangement shown in Fig. 1(b), where the system has separated out a number of clusters, but it has not yet reached the user's mental model. The user moves a few more images each time, iterating through the refinement process. After touching 20 images and asking for a total of 3 refinements, the user will see the clear clusters they were hoping to uncover, shown in Fig. 1(c). The total elapsed time, from the first touch to the final refinement, was approximately 90 seconds.

4 Conclusion

Most feature generation techniques, particularly for images, do not produce semantically relevant features. That is, features tend to be abstract numbers only meaningful in (often nonlinear) combination. Because users tend to see images based on semantic components (e.g., indoors, flowers, daytime) and not quantitative vectors (e.g., average contrast, SIFT, gradient histograms), their spatial arrangement of images will reflect the relationships the user identified from their domain knowledge. We present a tool to enable the user to produce those semantic arrangements and implicitly capture these data models as features for future image classification.

Future work will explore how many user-provided features essentially 'cover' the semantic space. Although there are potentially an infinite number of ways to arrange images on a canvas, users may only be able to distinguish a limited set of distinct arrangement strategies. How many different users would need to provide their own semantic features before new users would simply be referencing a superposition of existing features? If one relies on machine-generated features to cover most of the geometric or color features, would the number of semantic features needed be closer to 10 or 1000? This number is particularly relevant for evaluating the crowd-sourcing potential for this type of system.

Acknowledgments

The authors wish to thank Nathan Baker, Landon Sego and Michael Henry. This work is part of the Signature Discovery Initiative at Pacific Northwest National Laboratory. It was conducted under the Laboratory Directed Research and Development Program at PNNL, a multi-program national laboratory operated by Battelle for the U.S. Department of Energy.

References

[1] C. Andrews, A. Endert, and C. North. Space to Think: Large, High-Resolution Displays for Sensemaking. pages 55–64, 2010.

[2] E. T. Brown, A. Ottley, H. Zhao, Q. Lin, R. Souvenir, A. Endert, and R. Chang. Finding Waldo: Learning about Users from their Interactions. 2014.

[3] W. Dou, D. H. Jeong, F. Stukes, W. Ribarsky, H. R. Lipford, and R. Chang. Recovering Reasoning Processes from User Interactions. *IEEE Computer Graphics and Applications*, 29:52–61, 2009.

[4] A. Endert, L. Bradel, and C. North. Beyond control panels: Direct manipulation for visual analytics. *Computer Graphics and Applications, IEEE*, 33(4):6–13, 2013.

[5] P. Pirolli and S. Card. Sensemaking Processes of Intelligence Analysts and Possible Leverage Points as Identified Though Cognitive Task Analysis. *Proceedings of the 2005 International Conference on Intelligence Analysis, McLean, Virginia*, page 6, 2005.

[6] F. Shipman and C. Marshall. Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems. *Comput. Supported Coop. Work*, 8:333–352, 1999. 4.

(a) Initial Arrangement

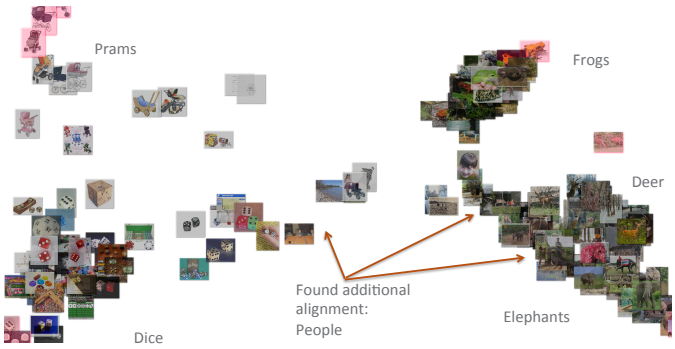

(b) First Refinement

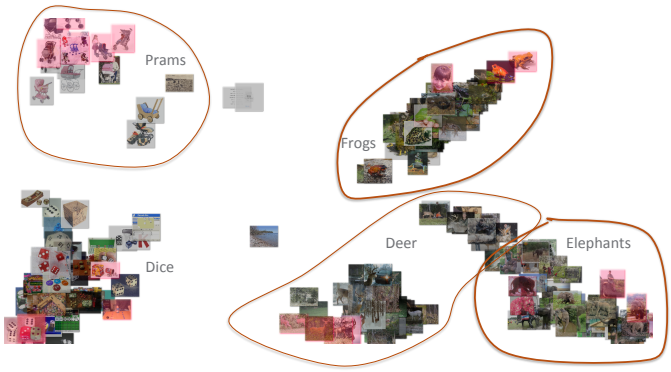

(c) Final Refinement, 20 images touched

Fig. 1. An example workflow. (a) The user initially positions 8 images (pink arrows). (b) After the first refinement, ActiveCanvas creates initial clusters. The user positions additional images to provide additional feedback to the canvas. (c) After touching a total of 20 images, and 3 refinements, clear clusters emerge. Although not readily apparent, the three images outside of clusters do not fit in any category.